EXHIBIT 2

IN THE UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF NEW YORK
ALBANY DIVISION

RENSSELAER POLYTECHNIC INSTITUTE AND DYNAMIC ADVANCES, LLC, Plaintiffs, v. APPLE INC., Defendant.	Civil Action No. 1:13-cv-00633-DEP JURY TRIAL DEMANDED

DECLARATION OF JAIME CARBONELL

I, Dr. Jaime Carbonell, declare as follows:

1. I have been retained as a technical expert on behalf of Rensselaer Polytechnic Institute and Dynamic Advances, LLC.

2. The following Declaration is based on my personal knowledge and all facts and statements contained herein are true and accurate to the best of my knowledge, information, and belief.

3. In preparing this Declaration, I have reviewed the following materials: U.S. Patent Number 7,177,798 (the "'798 Patent") and its file wrapper; and Apple Inc.'s Petition for Inter Partes Review of U.S. Patent 7,177,798 (the "IPR Petition") and all exhibits thereto.

4. It is my professional opinion that the IPR Petition is based on fundamentally flawed arguments and inferences. Through the IPR Petition, Apple argues to the U.S. Patent and Trademark Office that the claims of the '798 Patent are

invalid in light of certain prior art references (exhibits to the IPR Petition). Apple's invalidity arguments are premised on the prior art references disclosing every claim limitation of the claims of the '798 Patents. But it is clear that each of the references – alone and in combination – fail to disclose certain key limitations common to all claims of the '798 Patent. Thus, Apple's invalidity arguments facially and necessarily fail.

A. Qualifications

5. My qualifications for forming the opinions set forth in this report are summarized here and explained in more detail in my *curriculum vitae*, which is attached as **Exhibit A**. Additionally, in the past 5 years I have been an expert witness on a number of legal cases, primarily involving intellectual property in the areas of search engines, information retrieval, software, data mining and text mining.

6. I graduated from the Massachusetts Institute of Technology in 1975 with degrees in Physics and Mathematics. I went on to Yale University where I received a Masters Degree in Computer Science in 1976 and a Ph.D. in Computer Science in 1979.

7. In 1979, I became an Assistant Professor of Computer Science at Carnegie Mellon University. I was subsequently promoted to Associate Professor and then to Full Professor. Since 1995, I have been the Allen Newell Professor of Computer Science at Carnegie Mellon University. Since 1996, I have also been the Director of the Language Technologies Institute at Carnegie Mellon University. Last year, I was appointed "University Professor" at Carnegie Mellon University.

8. I have published over 300 technical and scientific articles, primarily in peer-reviewed journals and conferences in multiple computational fields, including: computer science, computational linguistics, natural language processing, machine

learning, databases, data mining, modeling, information retrieval, search engines, computational biology, machine translation, mathematical and statistical foundations, and integrated systems applications. These reflect my active lines of research over the past 35 years.

9. My research includes computational methods for analyzing text (typically called “natural language processing”) in order to organize it, retrieve it, summarize it, index it, parse it, and translate it. One of my papers describing my invention, Maximum Marginal Relevance for Retrieval and Summarization, is among the most highly cited (over 1,200 citations) in the Association for Computing Machinery’s Special Interest Group in Information Retrieval (ACM-SIGIR), the premier academic conference for search engines and related research. Among my other highly cited works is “Machine Learning: An Artificial Intelligence Approach” edited with Michalski and Mitchell (1,700 citations). I have researched mathematical approaches to analyze text, including statistical machine learning approaches over textual corpus, hand-built linguistic/heuristic methods, and combinations thereof. If asked, I can further discuss these areas of research and accomplishments. I have been a researcher in Natural Language Processing since my first paper on the subject in 1979. I have since published extensively in the field, as documented in my CV. From 1986 to 1996, I was the director of CMU’s Center for Machine Translation, which later evolved into the Language Technologies Institute that I still direct.

10. I teach courses and seminars in data mining, search engines, electronic commerce, machine learning, machine translation and aspects of computational biology at Carnegie Mellon University, mostly at the graduate level. I am also engaged in

designing distance-learning and learning-by-doing curricula, also at the graduate level. I advise Ph.D. and M.S. students in the above subject areas.

B. Person of Ordinary Skill in the Art

11. In my opinion, a person of ordinary skill in the art pertinent to the ’798 Patent at the time its application was filed would have a bachelor of science degree in computer science or a bachelor of science degree in engineering, mathematics, or physics with computer science coursework, and either graduate level coursework or 2-3 years of direct technical experience in database systems, and knowledge of natural language processing from courses or work experience.

C. Anticipation and Obviousness

12. Standard for Anticipation: I understand that to anticipate a claim, a reference must disclose each and every limitation of that claim, and that this should be assessed on a claim-by-claim basis. I also understand that anticipation can occur when an undisclosed limitation is literally missing, but is present because the prior art necessarily functions in accordance with, or includes, the undisclosed limitation.

13. I understand that for a reference to anticipate a patent claim, that reference must also enable one of ordinary skill in the art to make and use the full scope of the claimed invention without undue experimentation.

14. Standard for Obviousness: I understand that a combination of prior art references may render a claim obvious if, at the time of the invention, a person of ordinary skill in the art would have selected and combined those prior art elements in the normal course of research and development to yield the claimed invention. I understand in making the obviousness inquiry, one should consider the so-called *Graham* factors: the scope and content of the prior art; the differences between the

claimed inventions and the prior art; the level of ordinary skill in the art; and certain secondary considerations. I further understand the obviousness analysis is to be performed on a claim-by-claim basis. I understand that a person of ordinary skill in the art is a person of ordinary creativity, not an automaton. When there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp.

15. It is also my understanding that obviousness requires more than a mere showing that the prior art includes separate references covering each separate limitation in a claim under examination. I understand obviousness requires the additional showing that a person of ordinary skill at the time of the invention would have been motivated to select and combine those prior art elements in the normal course of research and development to yield the claimed invention. I also understand a fact-finder should be aware of the distortion caused by hindsight bias and must be cautious of arguments reliant upon ex post reasoning.

D. Natural Language Processing and Databases

16. Natural Language Processing (NLP) is a very large well-established discipline, almost since the start of computer science itself, and I have been an active researcher and teacher in NLP for over 30 years. NLP encompasses machine translation, text/web search, text mining, natural language database query, and more. Despite its long tradition, innovation was strong and steady in NLP in May 2001 (when the application for the ’798 Patent was filed) and continues at present. For instance, scientific conferences, each with hundreds of participants and dozens of presentations of new results abound: Association for Computational Linguistics (ACL), Empirical

Methods in Natural Language Processing (EMNLP), and International Conference on Computational Linguistics (COLING), to name a few. Patents in NLP continue to be issued in quantity; for instance, the USPTO lists 120 patents issued since January 1, 2013 with “natural language” in the abstract. The ’798 Patent itself discloses a large number of issued patents in the area as well as pertinent scientific publications.

17. NLP interfaces for databases was an active sub-area of NLP since the early days, continued to be so through the filing of the application for the ’798 Patent, and continues to this date, with steady innovation as evidenced by over 30 issued patents in this area in the past 2 years. The message should be clear: NLP for database query is both a fairly substantial area and an actively investigated one with steady innovation. There are many approaches to NLP in general, and to NLP for database query in particular. The simple presence of past work in NLP interfaces to databases is by no means evidence of lack of patentability; instead, the specifics of the patent claims in question, namely those of the ’798 Patent, must be analyzed in detail with respect to each alleged prior art.

E. The ’798 Patent

18. The ’798 Patent is directed squarely at natural language interfaces to databases (NLP for databases or NLP-DB for short). Relational databases were invented by Ted Codd (who received the Turing Award for his invention, the highest honor in Computer Science), and since then they became the dominant paradigm for databases, with new variants such as object-oriented and distributed databases. The standard means of querying databases is via SQL (Structured Query Language), which is in essence a special-purpose programming language. However many non-programmers, or programmers who do not know SQL, need to access databases

routinely, and therein lies the challenge – how to do so without specialized knowledge for programming SQL. One approach is to hire SQL programmers to act as intermediaries between the user and the database, but that is slow and expensive. Another approach is via a graphical user interface to SQL, but this limits query expressivity to predefined types of queries. The preferred approach relies on natural language interfaces, and that is the approach taken by the invention of the ’798 Patent.

19. The ’798 Patent distinguishes itself from prior art by addressing head-on the issue of linguistic coverage and query disambiguation through direct use of the database, metadata database, and reference dictionary (essentially a specification of the metadata database). The problem of “linguistic closure” – that is, coping with all possible linguistic variants of queries and elements therein is stated up front (’798 Patent col.2 ll.23-31): “A problem realized with many conventional natural language system designs is that these designs require exceedingly large collections of linguistic terms that users use, but still might not be able to assure successful closure of users’ queries. Because of design complexity and keyword data-base size, most systems are not practical to implement. A better approach to processing natural language inputs is therefore needed.”

20. The ’798 Patent goes on to critique as insufficient the prior approaches to NLP-DB, including implicitly those cited in Apple’s alleged prior art, for instance (’798 Patent col.5 ll.43-50): “Most previous research efforts on NLI [Natural Language Interfaces] have set out to find ways for the computer to understand the user’s articulation, following the more established tradition of Artificial Intelligence (AI). But their results have stopped significantly short of being truly natural. They have all endeavored to devise particular controls and limitations on the naturalness of input and

use these artifacts to assist interpreting queries into some standard database query languages."

21. Then the '798 Patent provides the key insight underlying the invention ('798 Patent col.6 ll.14-21): "One could argue that users are bound to refer, either directly or indirectly, to these known database objects (types or semantic models, instances or values, and operators) in their natural queries. If they do not use these database objects directly, they still have to use other words and phrases (henceforth known as "keywords") that correspond to these objects. Thus, the domain of interpretation is finite, compared to natural language processing in general." And the specification further states (*id*. col.6 ll.24-30): "The critical success factor of the last approach [the semantic model approach] depends clearly on the semantic model-dictionary employed, which must be powerful enough at least to span the range of possible usage . . . ." The specification goes on to describe the invention.

22. One key aspect of the invention is the set of ingredients used in the semantic model, as detailed in the '798 Patent (*id*. col.8 ll.51-59): "[A]ccording to one aspect of the invention, there are four layers of enterprise metadata (resources of search) considered; i.e., cases, keywords, information models, and database values. According to one aspect of the invention they are integrated in an extensible metadata representation method so that every resource item references all other related resources for query interpretation. A repository of metadata may be implemented as, for example, a reference dictionary." Further details follow. However, the above is the crux of the invention: instead of first performing opened-ended NLP – a daunting task not yet truly solved – use the database itself, together with the metadatabase semantic model, to restrict the possible words, word meanings, keywords, and semantic

interpretations to a manageable finite set. In other words, tailor the interface to the database and the metadata that describes it, rather than a two step process of general NLP followed by NLP-to-database-query mapping, since the first of these two traditional steps is much too difficult and error-prone. That is the essence of the invention, though, of course, the details provided in the specification are also material.

23. Claim 1 of the ’798 Patent recites:

A method for processing a natural language input provided by a user, the method comprising:

a. providing a natural language query input by the user;
b. performing, based on the input, without augmentation, a search of one or more language-based databases including at least one metadata database comprising at least one of a group of information types comprising:
    i. case information;
    ii. keywords;
    iii. information models; and
    iv. database values;
c. providing, through a user interface, a result of the search to the user;
d. identifying, for the one or more language-based databases, a finite number of database objects; and
e. determining a plurality of combinations of the finite number of database objects.

24. Element (b) of claim 1 should be interpreted as including a metadata database comprised of one or more groups of information types, each group containing

all four of: case information, keywords, information models, and database values. This reading is fully supported by the ’798 Patent: for example, the patent (col.8 ll.51-54) states: “there are four layers of enterprise metadata (resources of search) considered; i.e., cases, keywords, information models and database values, they are integrated in an extensible metadata representation method so that every resource[] item references all other related resources for query interpretation.” And Figure 2 contains all four information types in the “reference dictionary” (metadata) – as stated in the ’798 Patent (col.8 ll.57-59): “A repository of metadata may be implemented as, for example, a reference dictionary.” In other words, these four metadata elements are meant to function together, as a group, as claim 1 states. Claim 9 largely parallels claim 1 with respect to the group of information types, reciting “…at least one of a group of information comprising case information, keywords, information models, and database values.” Moreover a “group” cannot refer to a single element, and “database values” are not in themselves metadata (although they can be contained in a metadata database along with metadata). Hence the above reading is the one consistent with the specification, with the other claims, and with the normal meaning of “group” and of “metadata.”

25. Element (b) of claim 1 of the ’798 Patent, and the corresponding elements of claims 4, 9, and 14 are absent from each of the prior art references Apple relies on in its IPR Petition. None of Apple’s references disclose case information as metadata, and certainly none disclose the combination of all four types of metadata for NLP of database queries.

F. Validity of the ’798 Patent

26. In its IPR Petition, Apple argues that U.S. Patent 5,197,005 (inventors: Shwartz, et al., hereinafter “*Shwartz*”) anticipates claims 1-21 of the ’798 Patent. *Shwartz* does not anticipate the ’798 Patent claims at least because it fails to disclose the group of information types as claimed in the ’798 Patent.

27. *Shwartz* discloses a database retrieval system with “*A database independent, canonical internal meaning representation of a natural language query*” (*Shwartz*, Abstract (emphasis added)). The natural language query goes through a pipeline of steps divorced from the database or its associated metadata, and only at the end is it converted into an SQL query (*id*. Fig 1). A “semantic network” (i.e., a knowledge base or “internal meaning representation”) separate from any specific database is illustrated in Fig 7. In fact, Appendix A of *Shwartz* provides explicitly the Internal Meaning Representation Grammar, which does not use or rely upon any database or the claimed metadata database or reference dictionary or case information. Essentially, *Shwartz* represents an excellent example of the traditional paradigm for NLP-DB wherein the NLP is preformed independent of the database, contrary to the teachings of the ’798 Patent, wherein data and metadata play central roles in the NLP. Simply, *Shwartz* discloses technology quite different from the invention of the ’798 Patent.

28. In its IPR Petition, Apple also argues that Jurgen M. Janas’s article, *The Semantics-Based Natural Language Interface to Relational Databases*, in Cooperative Interfaces to Information Systems (Bolc & Jarke Eds.) (hereinafter “*Janas*”) anticipates claims 1-11, 13, 15, 16, and 18-21 of the ’798 Patent. *Janas* does not anticipate the

’798 Patent claims at least because it fails to disclose the group of information types as claimed in the ’798 Patent.

29. *Janas* discloses the NLP processing of queries into a formal query language (pp. 150-151), a simplified form of SQL or its equivalent, as a “translation process” (e.g., pp. 146, 154-156). Like *Shwartz*, as discussed above, and unlike the ’798 Patent, *Janas* does not rely on the database or the claimed metadata database to take part in the processing of the natural language query. Moreover, *Janas* makes no reference to past query processing or past disambiguation of queries, or any reuse of this dynamically recorded prior computation, and therefore fails to disclose “case information.” Instead, much of *Janas* is directed towards coping with linguistic phenomena such as conjunctions and ellipsis – important in their own right, but not pertinent to the ’798 Patent’s invention.

30. In its IPR Petition, Apple also argues that Beatrice Bouchou’s and Denis Maurel’s article, *Using Transducers in Natural Language Database Query*, in the 4th International Conference on Applications of Natural Language to Information Systems (hereinafter “*Bouchou*”) anticipates claims 1-6, 8-12, 14-16, and 20-21 of the ’798 Patent. *Bouchou* does not anticipate the ’798 Patent claims at least because it fails to disclose the group of information types as claimed in the ’798 Patent.

31. *Bouchou* proposes a very simple method for NLP of database queries, relying on finite-state machines called “transducers,” which essentially map one string into another (i.e., transducers map a word or sequence of words into another word or sequence of words). Transducers are less powerful than even the simplest grammars typically applied in NLP, namely context-free grammars. As such they have been used for simple tasks such as morphology, for instance generating the plural form of a word

from the singular by adding “s” or “es” at the end, adding “ed” at the end of a verb to change it to its past tense, or substituting a word with a synonymous one or equivalent one, the latter being illustrated in *Bouchou* (p. 9).

32. Moreover, *Bouchou* relies on linguistic methods (Sec 1.3 on pp. 8-9) stating “the use of linguistic operators rather than of logical operators,” albeit the simplistic linguistic methods mentioned above, and does not rely on semantic (logical) methods. *Bouchou* in fact teaches away from semantic methods, whether they use a general knowledge base or use the more specific and reliable semantics from the metadata database claimed in the ’798 Patent. This is further emphasized elsewhere in *Bouchou* (Sec 3.3 on p. 12). Nowhere does *Bouchou* disclose “case information” or “information models” and thus fails to disclose the group of information types as claimed in the ’798 Patent.

33. In its IPR Petition, Apple also argues that Sajul Dar, Gadi Entin, Shai Geva and Eran Palmon’s paper, *DTL’s DataSpot: Database Exploration Using Plain Language*, in the 24th VLDB Conference (hereinafter “*Dar*”) anticipates claims 9-10 and 13-16 of the ’798 Patent. *Dar* does not anticipate the ’798 Patent claims at least because it fails to disclose the group of information types as claimed in the ’798 Patent.

34. *Dar* describes “DTL’s DataSpot,” which is a “database publishing tool” and first converts databases into something Dar et al. call the “hyperbase” offline (abstract). “A hyperbase is a graph structure comprised of nodes, edges and node labels” (p. 646, col.1); it is no longer a database comprised of tables, rows, columns, entries. Then, the users may query the hyperbase “much in the same way as they use a Web search engine, such as Altavista” (p. 645, col.2); that is, by using key terms, such as “Mexico” or “Nancy’s sea food orders,” rather than natural language queries (p. 647, col.2). In

fact, *Dar* differentiates itself from systems containing natural language interfaces that translate queries into SQL (p. 649, col.1). Further, *Dar* completely fails to disclose case information, disclosing instead “thesaurus” and “heuristics” as powering its search for answers to queries. Simply, *Dar* does not disclose a method for natural language processing that includes reference to the group of information types claimed in the ’798 Patent.

35. In its IPR Petition, Apple also argues that U.S. Patent 6,584,464 (inventor: Warthen, hereinafter “*Warthen*”) anticipates claims 1-11 and 13-21 of the ’798 Patent. *Warthen* does not anticipate the ’798 Patent claims at least because it fails to disclose the group of information types as claimed in the ’798 Patent.

36. *Warthen* describes the AskJeeves search engine. Unlike other search engines, AskJeeves stores questions and answers, or more generally question templates (questions in canonical form such as “what is the weather in Boston”?) and answers (e.g., a URL to the weather.com page with Boston as the location). Questions in non-canonical form (such as “what’s Boston’s weather?” Or “For Boston, tell me the weather”) go through a multi-step process for conversion into canonical form, and then the answer is looked up in the question-answer mapping table and returned to the user (Fig 1(a) and 1:27-56). *Warthen* simply teaches a question-answer mapping table – the equivalent to a very large FAQ (Frequently-Asked Questions) list. Nothing new can ever be answered unless there is a canned answer in the FAQ (the question-answer mapping table). Hence, *Warthen* discloses a fairly impoverished model, and does not have the power of either a true search engine (such as Google, Bing, Alta Vista or Lycos), nor does it have the power of a natural language interface to a real database, where answers can be calculated by retrieving and *combining* multiple elements from

one or multiple tables. *Warthen* discloses nothing more than a big lookup list, like a phone book, albeit allowing a degree of flexibility via attempting to canonicalize the input. It does not disclose case information, metadata database information models, or the group of information types claimed in the '798 Patent.

37. Finally, in its IPR Petition, Apple also argues that claims 11, 12, 14, and 17 are obvious in light of David L. Waltz's article, *An English Language Question and Answering System for a Large Relational Database*, (hereinafter "*Waltz*"), in various combinations with *Janas*, *Warthen*, *Bouchou*, and *Dar*.

38. As with Apple's other references, *Waltz* fails to disclose at least the group of information types claimed in the '798 Patent. Indeed, Apple does not suggest that *Waltz* discloses such. Therefore the combinations of references including *Waltz* necessarily fail to disclose the group of information types as claimed in each of the claims of the '798 Patent.

G. Conclusion

39. Taken alone or in combination, the references that Apple relies on in its IPR Petition clearly fail to disclose certain key limitations common to all of the '798 Patent claims, including the case information and the group of information types as claimed in the patent. Apple's invalidity arguments in the IPR Petition, therefore, facially and necessarily fail.

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Executed on January 2, 2014,

J G Carbonell

Dr. Jaime Carbonell